arXiv:1609.09841v1 [math.NA] 30 Sep 2016

GPU Acceleration of Hermite Methods for the Simulation of Wave Propagation

Arturo Vargas, Jesse Chan, Thomas Hagstrom, Timothy Warburton

Abstract The Hermite methods of Goodrich, Hagstrom, and Lorenz (2006) use Hermite interpolation to construct high order numerical methods for hyperbolic initial value problems. The structure of the method has several favorable features for parallel computing. In this work, we propose algorithms that take advantage of the many-core architecture of Graphics Processing Units. The algorithm exploits the compact stencil of Hermite methods and uses data structures that allow for efficient data load and stores. Additionally the highly localized evolution operator of Hermite methods allows us to combine multi-stage time-stepping methods within the new algorithms incurring minimal accesses of global memory. Using a scalar linear wave equation, we study the algorithm by considering Hermite interpolation and evolution as individual kernels and alternatively combined them into a monolithic kernel. For both approaches we demonstrate strategies to increase performance. Our numerical experiments show that although a two kernel approach allows for better performance on the hardware, a monolithic kernel can offer a comparable time to solution with less global memory usage.

Arturo Vargas
Rice University, Houston TX USA e-mail: arturo.vargas@rice.edu.

Jesse Chan
Rice University, Houston TX USA.

Thomas Hagstrom
Southern Methodist University, Dallas TX USA.

Timothy Warburton
Virginia Tech, Blacksburg VA USA.

1 Introduction

Wave simulation is essential to many fields of study. For example, in geophysics the numerical solution to the acoustic wave equation is central to various imaging algorithms such as Reverse Time Migration [2] and Full Waveform Inversion [16]. In the context of electromagnetism, numerical simulations of Maxwell's equations are employed in the design of new products such as radars and antennae [14]. The need to resolve high frequency waves over long periods of time makes these simulations challenging. High order numerical methods can be more efficient than lower order methods for such simulations as they minimize dispersion and offer high convergence rates for smooth solutions [8].

The Hermite methods introduced by Goodrich et al. [5], are a class of cell based numerical methods which reconstruct a polynomial-based solution at each cell through Hermite interpolation. The methods can be constructed to achieve high order accuracy making them well suited for high frequency wave simulations. An advantageous feature of these methods is their high computation to communication ratio, making them ideal for parallel computing [3]. High performance implementations of Hermite methods have been carried out in [1, 6] for aero-acoustics and compressible flows in which numerical experiments demonstrated favorable results in terms of scalability on CPU-based clusters.

Recent trends in processor design has resulted in multi-core processors with wide single instruction multiple data (SIMD) vector units. Each SIMD group has access to a relatively small shared memory cache and each SIMD lane has a small number of fast registers. Typical GPUs are further equipped with large bandwidth, high latency, global shared memory storage. To achieve high performance on GPUs fine-grained parallelism must be exposed with minimal communication between computing units. Examples of numerical algorithms that have demonstrated utility of the GPU can be found in [10, 11, 9, 12]

Hermite methods were first implemented on a GPU by Dye in [4], wherein strategies for two-dimensional equations were presented. Building on the work of Dye, we introduce strategies for three-dimensional linear equations. In Section 2 we provide a brief overview of Hermite methods. Section 3 introduces our strategies for tailoring Hermite methods onto the GPU and lastly Section 4 studies performance with respect to the GPU hardware.

2 Overview of Hermite Methods

To highlight key concepts of Hermite methods, we consider the three-dimensional advection equation,

$$\frac{\partial u}{\partial t} = \frac{\partial u}{\partial x_1} + \frac{\partial u}{\partial x_2} + \frac{\partial u}{\partial x_3}. \tag{1}$$

Hermite methods represent the solution of an initial value boundary problem on a grid constructed through tensor products of one-dimensional grids. We denote the m^{th} node for the k^{th} dimension as x_{k,m_k} and for simplicity consider periodic grids. The degrees of freedom of the method, at time step $t_n = t_0 + n\Delta t$, are represented over each node in the form of the tensor product of the function value and first N (scaled) derivatives in each dimension,

$$p^i_{m_1,m_2,m_3}(t_n) \approx \frac{h^{|i|}}{i!} D^i u\left(x_{1,m_1}, x_{2,m_2}, x_{3,m_3}, t_n\right). \tag{2}$$

Here h denotes the spacing between the nodes, D denotes the derivative operator, and $i = (i_1, i_2, i_3)$ denotes the multi-index with i_j ranging from 0 to N.

To represent the solution on each cell of the grid a staggered (dual) grid is introduced. The cell midpoints of the primary grid make up the dual grid. The dual grid facilitates the construction of tensor polynomials (Hermite interpolants)

$$Rp_{m_1+\frac{1}{2},\ldots,m_3+\frac{1}{2}} = \sum_{j_1=0}^{2N+1} \cdots \sum_{j_3=0}^{2N+1} b_{j_1,\ldots,j_3} \left(\frac{x_1 - x_{1,m_1+\frac{1}{2}}}{h_{x_1}}\right)^{j_1} \ldots \left(\frac{x_3 - x_{3,m_3+\frac{1}{2}}}{h_{x_3}}\right)^{j_3}, \tag{3}$$

which interpolate the function value and derivatives at each cell's vertices. The coefficients of the tensor product polynomial are the approximation of the function value and the derivatives at the midpoint of the cell.

Evolution from t_n to $t_{n+\frac{1}{2}}$ is carried out independently on each cell by the use of a q-order temporal Taylor series expansion centered at t_n of the tensor product polynomial

$$TRp = \sum_{j_1=0}^{2N+1} \ldots \sum_{j_3=0}^{2N+1} \sum_{s=0}^{q} b_{j_1,\ldots,j_3,s} \left(\frac{x_1 - x_{1,m_1+\frac{1}{2}}}{h_{x_1}}\right)^{j_1} \ldots \left(\frac{x_3 - x_{3,m_3+\frac{1}{2}}}{h_{x_3}}\right)^{j_3} \left(\frac{t - t_n}{\Delta t}\right)^{s}. \tag{4}$$

The scalar Δt denotes the size of a full time step. For $s = 0$ the coefficients of Equation 4 are simply the coefficients from the Hermite interpolant (Equation 3). The time-stepping scheme of Hermite methods, Hermite-Taylor, expresses the values of unknown coefficients in terms of known coefficients by applying the Cauchy-Kowalweski recurssion to the PDE. For brevity, we omit the derivation and provide the resulting recursion for the three-dimensional advection equation,

$$b_{j_1,j_2,j_3,s+1} = \frac{j_1+1}{s+1}\frac{\Delta t}{h_{x_1}} b_{j_1+1,j_2,j_3,s} + \frac{j_2+1}{s+1}\frac{\Delta t}{h_{x_2}} b_{j_1,j_2+1,j_3,s} + \frac{j_3+1}{s+1}\frac{\Delta t}{h_{x_3}} b_{j_1,j_2,j_3+1,s}. \tag{5}$$

With the determined coefficients, the function value and derivatives for the midpoint are computed by evaluating the series at $t_{n+1/2}$. To complete a full time step the process is repeated on the dual grid to approximate the solution on the primary grid.

Hermite methods converge at a rate of $O(h^{2N+1})$ for smooth solutions, and are stable as long as the waves do not propagate from the cell boundaries to the cell center in a half step. A significant feature of the method's stability is that the result is independent of order. We refer the reader to [5, 7, 15] for further details on the methods.

By exploiting the compact stencil (the vertices of a cell) and local evolution of the method, we expose opportunities for parallelism. At a coarse level the polynomial reconstruction and evolution can be performed independently for each cell. At a finer level many of the operations can be carried out as one-dimensional matrix-vector multiplications. Because of the two levels of parallelism we demonstrate how the method can be mapped on to the many-core architecture of the GPU.

3 Implementing Hermite Methods on Graphics Processing Units

To simplify the performance analysis, we first implement the interpolation and evolution procedure as separate kernels. The drawback of this approach is the additional temporary memory required to store the reconstructed polynomial and additional memory transfers. In an effort to minimize global memory usage we implement a monolithic kernel performing both the interpolation and evolution.

Computation on the GPU is performed on a predefined grid of compute units. Following NVIDIA's nomenclature each unit of the grid is referred to as a thread. Threads are grouped to form thread blocks. The hardware provides a similar hierarchy for memory. Threads are provided with a small amount of exclusive memory, threads in a thread block share block exclusive memory (shared memory), and lastly the entire compute grid shares global memory. Moving data between the CPU and GPU is accomplished through the use of global memory which acts as a general buffer. We refer the reader to [13] for a detailed overview on GPU computing. All numerical experiments in this work are written in the Open Concurrent Compute Abstraction (OCCA) API [10] allowing for portability across hardware. Numerical experiments are carried out using an NVIDIA GTX 980 GPU in single precision using OCCA generated Compute Unified Device Architecture (CUDA) code. The hardware has theoretical peak bandwidth of 224 GB/sec and floating point performance of 4,612 GFLOP/sec.

3.1 Hermite Interpolation on the GPU

The fundamental data structure used throughout our implementation is the tensor. For example the tensor

$$\mathbf{u}[m_3][m_2][m_1][n_3][n_2][n_1],$$

is used to store the function value and derivatives at each node of a three-dimensional grid. The three innermost indices correspond to a grid point on the grid and the outermost indices catalog the corresponding tensor product of function value and derivatives.

In three dimensions, polynomial reconstruction at a node on the dual grid, is accomplished by interpolating the function value and derivatives from vertices of

the encapsulating cell. This requires reading $(N+1)^3$ degrees of freedom per vertex, for a total of eight vertices in three dimensions.

To facilitate the interpolation procedure a one-dimensional Hermite interpolation operator, $\mathbf{H}$ (see [15] for details on construction), is pre-computed enabling dimension-by-dimension reconstruction of the polynomial. In this kind of reconstruction, the degrees of freedom of the encapsulating cell are stored in a local rank 3 tensor, $\mathbf{u_{loc}}$. The one-dimensional operator, $\mathbf{H}$, is then applied to the degrees of freedom of nodes parallel to the x_1 dimension as a series of matrix-vector multiplications. Next, the operator is applied to the degrees of freedom of nodes parallel to the x_2 dimension, and lastly to the degrees of freedom of nodes parallel to the x_3 dimension. For clarity we define $\mathbf{H}_{x_1}$, $\mathbf{H}_{x_2}$, and $\mathbf{H}_{x_3}$ as operators to be applied in the x_1, x_2, and x_3 dimensions respectively. Algorithm 1 presents the application of the interpolation operator to nodes parallel to the x_1 dimension using nested for loops. Applying the operator in the x_2, and x_3 dimensions is performed analogously. The complete reconstruction procedure for a single polynomial is listed as Algorithm 2.

Algorithm 1 Polynomial reconstruction in the x_1 dimension

```
1:  procedure RECONSTRUCTIONINx1(H_x1, u_loc, Ru)
2:      for tz=0:2N+1 do
3:          for ty=0:2N+1 do
4:              for tx=0:2N+1 do
5:                  c=0
6:                  for k=0:2N+1 do
7:                      c += H_x1[tx][k] u_loc[tz][ty][k]
8:                  end for
9:                  Ru[tz][ty][tx]=c;
10:             end for
11:         end for
12:     end for
13: end procedure
```

Algorithm 2 Polynomial reconstruction

```
1: procedure POLYNOMIALRECONSTRUCTION(H_x1, H_x2, H_x3, u_loc, Ru)
2:     Ru = H_x1 u_loc
3:     u_loc = H_x2 Ru
4:     Ru = H_x3 u_loc
5: end procedure
```

Our GPU implementation exposes two levels of parallelism: coarse parallelism, in which threads in a block collectively reconstructs polynomials, and fine-grain parallelism in which threads apply the interpolation operator as a series of matrix-vector multiplications. The reconstruction is carried out locally by moving the necessary degrees of freedom to shared memory.

To minimize and reuse global memory reads we apply a similar register rolling technique as used in Finite Difference Time Domain methods [11]. Hermite methods can mimic this technique by having a block of threads reuse a subset of shared memory. This is accomplished by setting up a two-dimensional grid of thread blocks. A single block of threads moves the bottom four vertices of a cell to shared memory. The block of threads then applies the interpolation operators $\mathbf{H}_{x_1}$ and $\mathbf{H}_{x_2}$. As it progresses along the x_3-dimension it stores the next four vertices of the cell in shared memory and applies $\mathbf{H}_{x_1}$ and $\mathbf{H}_{x_2}$ to the newly added degrees of freedom. As there are now degrees of freedom for eight vertices, the $\mathbf{H}_{x_3}$ operator is then applied to the degrees of freedom parallel to the x_3 dimension and the result is stored in a rank 6 tensor similar to the initial degrees of freedom. The block of threads then shifts forward to the next set of four nodes and repeats the polynomial reconstruction.

We add a tunable parameter: the number of polynomials reconstructed along the x_1-dimension per block of threads. This can further reduce the total amount of global memory reads as neighboring cells share nodes on the interface. Figure 1 reports the performance for the polynomial reconstruction kernel under a naive implementation, with no reuse of existing memory reads, and the optimized kernel with reuse of global memory reads. Figure 2 visualizes the relationship between global throughput and bandwidth. As we increased data reuse we observed higher bandwidth. Additionally for the reconstruction of order 3, and 5 polynomials, $N = 1,2$ respectively, there was a reduction in throughput in the optimized kernel suggesting caches are being exploited.

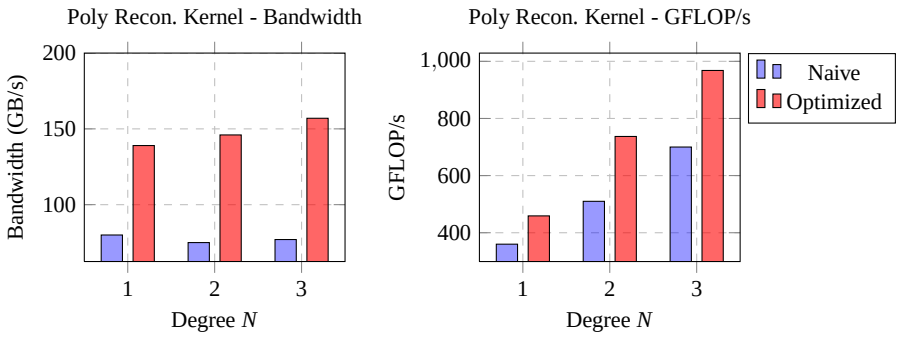

Fig. 1: Performance of the interpolation kernel. The optimized kernel assigned the construction of 16, 10, and 4 interpolants to each block threads for $N = 1,2,3$, reconstructing order 3, 5, and 7 polynomials respectively.

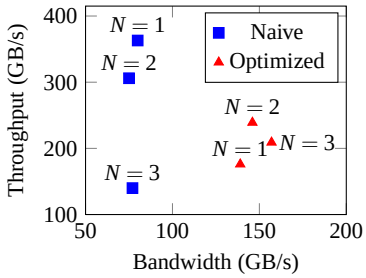

Fig. 2: Plotting bandwidth against throughput shows that assigning more cells per block reduces global throughput for orders $N = 1, 2$, which reconstruct polynomials of orders 3, and 5, while increasing global bandwidth.

3.2 Hermite-Taylor Methods on the GPU

With the polynomial reconstruction procedure described in the previous section completed we may now advance the solution using the Hermite-Taylor method. For each reconstructed polynomial, the procedure can be performed locally using a rank 3 tensor to store the coefficients,

$$(\mathbf{Ru})[n_3][n_2][n_1],$$

where n_3, n_2, n_1 range from $0, \cdots, 2N+1$, corresponding to the order of spatial derivative in each spatial dimension. Differentiating the reconstructed polynomial with respect to a spatial dimension is achieved by applying the following derivative matrix

$$\mathbf{D}_{ij} = \begin{cases} \frac{i+1}{h} & , \quad j = i+1 \\ 0 & , \quad \text{otherwise}, \end{cases} \qquad 0 \leq i, j \leq 2N+2,$$

to the reconstructed polynomial. For convenience $\mathbf{D}_{x_1}$ will represent an operator to be applied to the reconstructed polynomial with respect to the x_1 dimension. In a similar manner the operators $\mathbf{D}_{x_2}$, and $\mathbf{D}_{x_3}$ will represent an operator to be applied to the reconstructed polynomial with respect to the x_2, and x_3 dimensions. For example application of the derivative matrix along the x_1 dimension to the reconstructed polynomial is illustrated in Algorithm 3 using nested for loops. Differentiating the reconstructed polynomial in the remaining dimensions is accomplished analogously. With the compact notation the Hermite-Taylor algorithm can be reduced to a q-stage loop as listed in Algorithm 4. Carrying out $q = d(2N+1)$ stages, in d dimensions, allows for the largest possible time step. Taking $q < d(2N+1)$ corresponds to a lower order temporal approximation and may require a smaller time-step in order to maintain expected order of convergence and stability. Similar to the polynomial reconstruction kernel, we consider two levels of parallelism: a coarse level in which each block of threads carries out the Hermite-Taylor scheme for a number of cells and a fine-grained level in which threads to carry out the computation. Numerical

Algorithm 3 Differentiation in the x_1-dimension

```
1: procedure DIFFERENTIATIONIN x1(D_x1, Ru, Ru_x)
2:     for tz = 0, 2N+1 do
3:         for ty = 0, 2N+1 do
4:             for tx = 0, 2N+1 do
5:                 if tx < 2N+1 then
6:                     p_x1 = ((tx+1)/h_x) Ru[tz][ty][tx+1]
7:                 else
8:                     p_x1 = 0
9:                 end if
10:                Ru_x[tz][ty][tx] = p_x1
11:            end for
12:        end for
13:    end for
14: end procedure
```

Algorithm 4 Hermite-Taylor evolution

```
1: procedure HERMITE-TAYLOREVOLUTION(D_x1, D_x2, D_x3, Ru)
2:     ŵ = Ru
3:     for k = q, q−1, ..., 1 do
4:         ŵ = Ru + (Δt/k)(D_x1 ŵ + D_x2 ŵ + D_x3 ŵ)
5:     end for
6:     Ru = ŵ
7: end procedure
```

experiments demonstrated that increasing the number of stages, q, in the scheme increases computational intensity. Peak performances were observed when assigning a block of threads to evolve the solution at 16, 10, and 2 cells for orders N=1, 2, and 3 respectively. Performance results are reported in Figure 3.

3.3 A Monolithic Kernel

A two kernel approach allows for fine tuning of each individual procedure at the cost of storing the coefficients for the reconstructed polynomial. In the interest of minimizing global storage we combine the polynomial reconstruction and evolution procedures to a single monolithic kernel. We repeat previous experiments carried out in Section 3.2 and observe the relationship between number of stages in the Hermite-Taylor scheme and performance. Figure 4 reports the performance for the monolithic kernels.

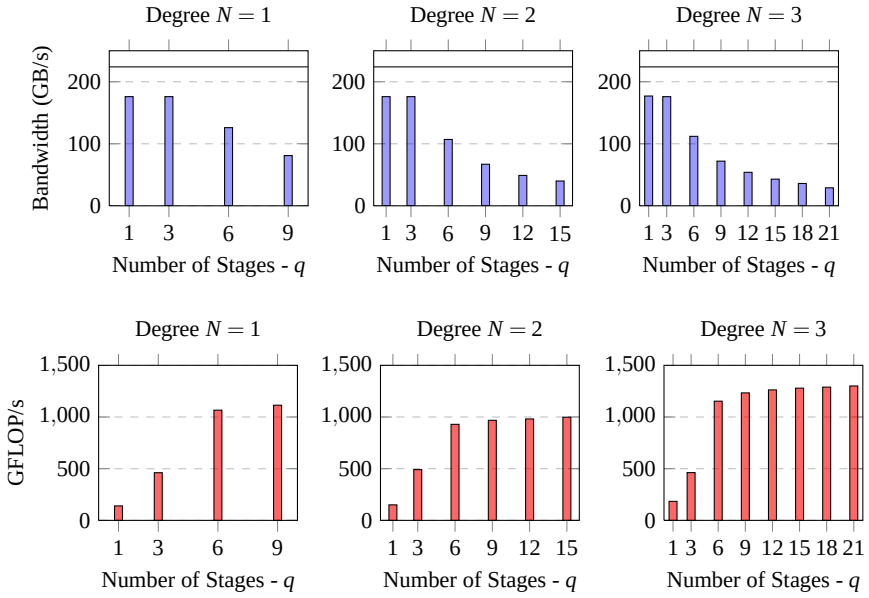

Fig. 3: Performance of the Hermite-Taylor kernel, the kernel assigns the evolution at 16, 10, and 2 cells per block of threads. An N^{th} degree method reconstructs local order $2N+1$ polynomials. As the order of the temporal expansion increases the kernel becomes more compute intensive. The solid line on the bandwidth plot denotes the peak theoretical bandwidth of the device.

4 Roofline Analysis and Time to Solution

The Roofline model relates flops, bandwidth, and hardware [17]. It provides an upper bound on the rate of floating point operations based on the arithmetic intensity of a given kernel. Arithmetic intensity is defined as:

$$\text{arithmetic intensity} = \frac{\text{FLOPs performed}}{\text{bytes loaded}}.$$

Pairing the arithmetic intensity and the physical capabilities of the hardware allows the roofine model to present a theoretical ceiling on performance for a given kernel. Theoretical achievable performance is defined as,

$$\min(\text{arithmetic intensity} \times \text{peak bandwidth}, \text{peak GFLOP/s}).$$

Figure 5 profiles the Hermite kernels in this work with respect to the Roofline model and reports the computational efficiency. The Hermite-Taylor and monolithic kernels are profiled using a $q = d(2N+1)$ stage loop. Typically there are two types of computational bottle necks, bandwidth or compute. Kernels which are bandwidth limited are constrained by a device's ability to read and write to global memory.

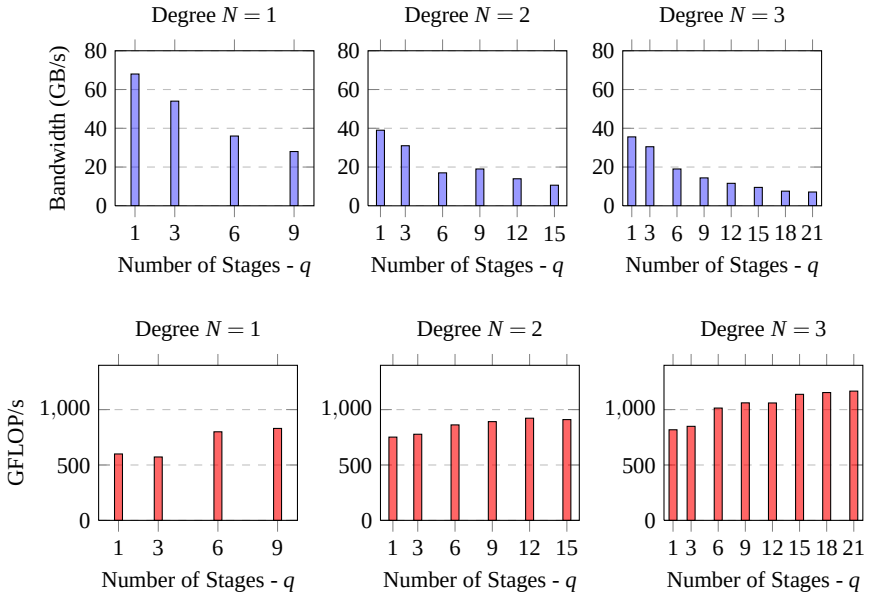

Fig. 4: Performance of the monolithic kernel. An N^{th} degree method reconstructs local order $2N+1$ polynomials. As the order of the temporal expansion increases the kernel becomes more compute intensive. Peak performances were found when assigning 12, 10, 2 cells per block of threads for orders $N = 1,2,3$ respectively.

Compute bound kernels are limited by the device's ability to perform floating point operations. The Roofline model places kernels limited by bandwidth on the bottom left while compute bound kernels are found on the top right. We observe that our kernels have a higher compute intensity and are closer to being compute bound. This is largely due to the high number of stages in the Hermite-Taylor scheme. Reducing the number of stages reduces the floating point intensity. Noticeably the interpolation and evolution kernels achieve a higher hardware efficiency.

Although separate kernels for the interpolation and evolution lead to better hardware efficiency, computational experiments have demonstrated that both approaches lead to comparable times to solution. The monolithic kernel has the advantage of requiring less reads/writes to global memory in comparison to the two kernel approach. Table 1 reports a comparison of time to solution for the advection equation on a fixed grid with 150 grid points in each dimension propagated for 200 timesteps. The caveat is that the local variables must be able to fit in shared memory when using a monolithic kernel. We carried out similar experiments for the acoustic wave equations and have found that peak performances were found by reducing the number of cells per block relative to the advection equation. The additional variables increases the use of hardware resources.

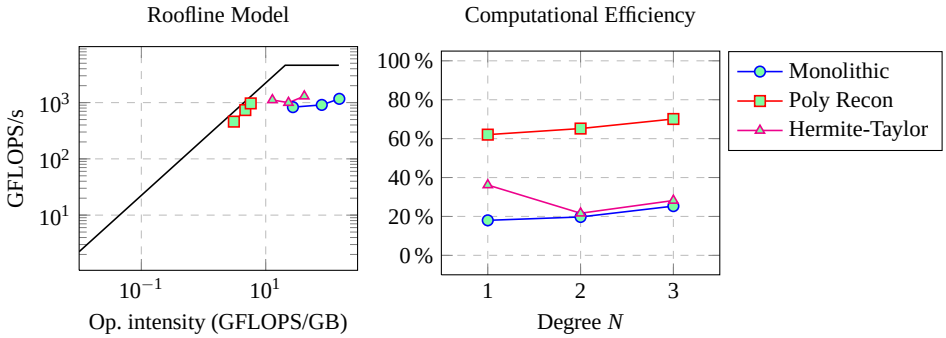

Fig. 5: Roofline performance analysis for the various Hermite method kernels.

	$N=1$	$N=2$	$N=3$
Advection: Monolithic Kernel	2.09 (sec)	13.77 (sec)	36.17 (sec)
Advection: Two Kernels	2.42 (sec)	13.77 (sec)	37.23 (sec)

Table 1: Comparison of time to solution. The initial condition is propagated for 200 time-steps on a fixed grid of 150 grid points in each dimension. A degree N Hermite method converges at a rate of $O(h^{2N+1})$.

5 Conclusion

This work examines the use of a GPU as a kernel accelerator for Hermite methods. Hermite methods consist of two main components, the reconstruction of a polynomial of order $2N+1$ and evolution via a space-time expansion. We presented two strategies in which to exploit the many-core architecture of the GPU. The first considered separate kernels for the polynomial reconstruction and evolution while the second considered a monolithic kernel. We demonstrated that separate kernels for the polynomial reconstruction and evolution make better use of the hardware capabilities but the fewer global memory read/writes of a single monolithic kernel enables for a comparable time to solution with less global memory usage. Future work will examine optimization strategies in the case of spatially varying coefficients and the employment of multiple GPUs.

Acknowledgements TH was supported in part by NSF Grant DMS-1418871. TW and JC were supported in part by NSF Grant DMS-1216674. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the National Science Foundation.

References

1. D. Appelö, M.Inkman, T. Hagstrom, and T. Colonius Recent progress on Hermite methods in aeroacoustics. In 17th AIAA/CEAS Aeroacoustics Conference. AIAA, 2011.
2. E. Baysal, D. D. Kosloff, and J.W. Sherwood. Reverse time migration. Geophysics, 48:1514-1524, 1983.
3. X. Chen. Numerical and Analytical Studies of Electromagnetic Waves: Hermite Methods, Supercontinuum Generation, and Multiple Poles in the SEM, Doctoral Thesis, University of New Mexico, 2012.
4. E. T. Dye. Performance Analysis and Optimization of Hermite Methods on NVIDIA GPUs Using CUDA, Master Thesis, The University of New Mexico, 2015.
5. J. Goodrich, T. Hagstrom, and J. Lorenz. Hermite methods for hyperbolic initial-boundary value problems. Math. Comp., 75:595630, 2006.
6. T. Hagstrom, D. Appelö. Experiments with Hermite methods for simulating compressible flows: Runge-Kutta time-stepping and absorbing layers. In 13th AIAA/CEAS Aeroacoustics Conference. AIAA, 2007.
7. T. Hagstrom, D. Appelö, 2015. Solving PDEs with Hermite Interpolation. In Spectral and High Order Methods for Partial Differential Equations ICOSAHOM 2014 (pp. 31-49). Springer International Publishing, 2014.
8. J. S. Hesthaven, and T. Warburton. Nodal discontinuous Galerkin methods: algorithms, analysis, and applications. Springer Science & Business Media, 2007.
9. A. Klöckner, T. Warburton, J. Bridge, and J. S. Hesthaven. Nodal discontinuous Galerkin methods on graphics processors. Journal of Computational Physics, 228:7863-7882, 2009.
10. D. Medina. OKL: A Unified Language for Parallel Architectures, Doctoral Thesis, Rice University, 2015.
11. P. Micikevicius. 3D finite difference computation on GPUs using CUDA. In Proceedings of 2nd workshop on general purpose processing on graphics processing units (pp. 79-84), ACM, 2009.
12. A. Modave, A. St-Cyr, and T. Warburton. GPU performance analysis of a nodal discontinuous Galerkin method for acoustic and elastic models. Computers & Geosciences, 91:64-76, 2006.
13. J. Sanders, and E. Kandrot. CUDA by example: an introduction to general-purpose GPU programming. Addison-Wesley Professional, 2010.
14. A. Taove, and S. C. Hagness. Computational electrodynamics: the finite-difference time-domain method. Norwood, 2nd Edition, MA: Artech House, 1995.
15. A. Vargas, J. Chan, T. Hagstrom, and T. Warburton. Variations on Hermite methods for wave propagation. arXiv preprint arXiv:1509.08012, 2015.
16. J. Virieux, and S. Operto, An overview of full-waveform inversion in exploration geophysics. Geophysics, 74:WCC1-WCC2, 2009.
17. S. Williams, A. Waterman, and D. Patterson. Roofline: an insightful visual performance model for multicore architectures. Communications of the ACM, 52:65-76, 2009.